

Amendments to the Claims:

This listing of claims will replace all prior versions, and listings, of claims in the application:

Listing of Claims:

1. - 32. (canceled)

33. (currently amended) An apparatus ~~A data repository node~~, comprising
an interface operative to communicate with client nodes and at least one other distributed data repository node over a computer network, and
a mapping module comprising a content map stored in a memory, the content map comprising one or more content map entries, each content map entry comprising a unique identifier and one or more record chunks associated with the unique identifier, each of the record chunks comprising a binary data object, and at least one index map stored in the memory, the at least one index map comprising one or more index map entries, each index map entry comprising a unique identifier corresponding to one or more record chunks maintained in the content map and one or more record attribute values associated with corresponding ones of the binary data objects of the record chunks; wherein the mapping module is operative to

receive a request to insert a record from a first client node;

generate a unique identifier in response to the record insertion request;

transmit the unique identifier to the first client node;

receive an insertion message including the unique identifier and at least one record attribute value;

store the at least one record attribute value associated with the insertion message in a corresponding index map in association with the unique identifier;

receive record chunks of a data stream corresponding to the unique identifier

from the first client node;

store the record chunks in the content map in association with the unique identifier;

receive a query from a second client node;

access one or more of the at least one index map to identify a unique identifier corresponding to one or more record chunks that, ~~wherein the record~~ satisfies the query;

provide, to the second client node, the record chunks associated with the unique identifier ~~corresponding to the record~~; and

stream additional record chunks associated with the unique identifier of the ~~data stream~~ to the second client node as they are received from the first client node.

34. (previously amended) The data repository node of claim 33 wherein the mapping module is further operative to

synchronize the record attribute values in the at least one index map with record attribute values of at least one index map maintained by the at least one other distributed data repository node.

35. (currently amended) The data repository node of claim 34 wherein the mapping module is further operative to transmit the record chunks to at least one other data repository node for replication.

36. (currently amended) A distributed data repository system, comprising

at least two distributed repository nodes, each distributed repository node comprising

a content map stored in a memory, the content map containing at least one message payload stored in association with a message payload identifier, wherein the at

least one message payload comprises a binary data object, and

at least one index map stored in the memory, the at least one index map
containing at least one content attribute value associated with a corresponding binary data
object and a corresponding message payload identifier;

wherein each distributed repository node is operative to

receive a request to insert a record from a first client node;

generate a first unique message payload identifier in response to the record
insertion request;

transmit the first unique message payload identifier to the first client node;

receive, from the first client node, an insertion message including the first
unique message payload identifier and at least one ~~record~~ content attribute value;

store the at least one record attribute value associated with the insertion
message in a corresponding index map in association with the first unique message
payload identifier;

receive, from the first client node, ~~record chunks~~ message payloads of a data
stream corresponding to the first unique message payload identifier from the client node;

store the record chunks in the content map in association with the first
unique message payload identifier;

receive a query from a second client node;

access one or more of the at least one index map to identify a unique message
payload identifier corresponding to one or more message payloads that, wherein the
~~record~~ satisfies the query;

provide, to the second client node, the ~~record chunks~~ message payloads
associated with the identified unique message payload identifier ~~corresponding to the~~
~~record~~;

stream additional ~~record chunks~~ message payloads associated with the
identified unique message payload identifier of the data stream to the second client node

as they are received from a third ~~the first~~ client node;

synchronize the ~~record~~ content attribute values in the at least one index map with ~~record~~ content attribute values of at least one index map maintained by the at least one other distributed data repository nodes.

37. (currently amended) The data repository ~~node~~ system of claim 36 wherein each distributed data repository node is further operative to transmit the ~~record-chunks~~ message payloads to at least one other data repository node for replication.

38. (currently amended) The data repository ~~node~~ system of claim 36 wherein each distributed data repository node is further operative to request and receive from at least one other distributed data repository node ~~record-chunks~~ message payloads that match a query received from a client node.

39. (currently amended) The apparatus ~~data repository node~~ of claim 33 wherein each distributed data repository node is further operative to transmit notifications to other distributed data repository node to reserve ~~reserving~~ the unique identifier.

40. (currently amended) The data repository ~~system~~ of claim 36 wherein each distributed data repository node is further operative to transmit notifications to other distributed data repository node to reserve ~~reserving~~ the first unique message payload identifier.

41. (new) The apparatus of claim 33 wherein the binary data objects are Binary Large Objects (BLOBs).

42. (new) The apparatus of claim 33 wherein the binary data objects are media content objects.

43. (new) The data repository system of claim 36 wherein the identified unique message payload identifier is the first unique message payload identifier and the third client node is the first client node.